

Role of homeostasis in learning sparse representations

Laurent U. Perrinet*

Institut de Neurosciences Cognitives de la Méditerranée (INCM)
CNRS / University of Provence
13402 Marseille Cedex 20, France
e-mail: Laurent.Perrinet@incm.cnrs-mrs.fr

Abstract

Neurons in the input layer of primary visual cortex in primates develop edge-like receptive fields. One approach to understanding the emergence of this response is to state that neural activity has to efficiently represent sensory data with respect to the statistics of natural scenes. Furthermore, it is believed that such an efficient coding is achieved using a competition across neurons so as to generate a sparse representation, that is, where a relatively small number of neurons are simultaneously active. Indeed, different models of sparse coding coupled with Hebbian learning and homeostasis have been proposed that successfully match the observed emergent response. However, the specific role of homeostasis in learning such sparse representations is still largely unknown. By quantitatively assessing the efficiency of the neural representation during learning, we derive a cooperative homeostasis mechanism which optimally tunes the competition between neurons within the sparse coding algorithm. We apply this homeostasis while learning small patches taken from natural images and compare its efficiency with state-of-the-art algorithms. Results show that while different sparse coding algorithms give similar coding results, the homeostasis provides an optimal balance for the representation of natural images within the population of neurons. Competition in sparse coding is optimized when it is fair: By contributing to optimize statistical competition across neurons, homeostasis is crucial in providing a more efficient solution to the emergence of independent components.

Keywords

Neural population coding, Unsupervised learning, Statistics of natural images, Simple cell receptive fields, Sparse Hebbian Learning, Adaptive Matching Pursuit, Cooperative Homeostasis, Competition-Optimized Matching Pursuit

*<http://www.incм.cnrs-mrs.fr/LaurentPerrinet/Publications/Perrinet10shl>.

1 Introduction

The central nervous system is a dynamical, adaptive organ which constantly evolves to provide optimal decisions for interacting with the environment. The early visual pathways provide a powerful system for probing and modeling these mechanisms. For instance, it is observed that edge-like receptive fields emerge in simple cell neurons from the input layer of the primary visual cortex of primates (Chapman & Stryker, 1992). The development of cortical cell orientation tuning is an activity-dependent process but it is still largely unknown how neural computations implement this type of unsupervised learning mechanisms. A popular view is that such a population of neurons operates so that relevant sensory information from the retino-thalamic pathway is transformed (or “coded”) efficiently. Such efficient representation will allow decisions to be taken optimally in higher-level layers or areas (Atick, 1992; Barlow, 2001). It is believed that this is achieved through lateral interactions which remove redundancies in the neural representation, that is, when the representation is *sparse* (Olshausen & Field, 1996). A representation is sparse when each input signal is associated with a relatively small number of simultaneously activated neurons within the population. For instance, orientation selectivity of simple cells is sharper than the selectivity that would be predicted by linear filtering. As a consequence, representation in the orientation domain is sparse and allows higher processing stages to better segregate edges in the image (Field, 1994). Sparse representations are observed prominently with cortical response to natural stimuli, that is, to behaviorally relevant sensory inputs (Baudot et al., 2004; DeWeese et al., 2003; Vinje & Gallant, 2000). This reflects the fact that, at the learning time scale, coding is optimized relative to the statistics of natural scenes. The emergence of edge-like simple cell receptive fields in the input layer of the primary visual cortex of primates may thus be considered as a coupled coding and learning optimization problem: At the coding time scale, the sparseness of the representation is optimized for any given input while at the learning time scale, synaptic weights are tuned to achieve on average optimal representation efficiency over natural scenes.

Most of existing models of unsupervised learning aim at optimizing a cost defined on prior assumptions on representation’s sparseness. These sparse learning algorithms have been applied both for images (Doi et al., 2007; Fyfe & Baddeley, 1995; Olshausen & Field, 1996; Perrinet, 2004; Rehn & Sommer, 2007; Zibulevsky & Pearlmutter, 2001) and sounds (Lewicki & Sejnowski, 2000; Smith & Lewicki, 2006). For instance, learning is accomplished in SPARSENET (Olshausen & Field, 1996) on patches taken from natural images as a sequence of coding and learning steps. First, sparse coding is achieved using a gradient descent over a convex cost derived from a sparse prior probability distribution function of the representation. At this step of the learning, it is performed using the current state of the “dictionary” of receptive fields. Then, knowing this sparse solution, learning is defined as slowly changing the dictionary using Hebbian learning. In general, the parameterization of the prior has major impacts on results of the sparse coding and thus on the emergence of edge-like receptive fields and requires proper

tuning. In fact, the definition of the prior corresponds to an objective sparseness and does not always fit to the observed probability distribution function of the coefficients. In particular, this could be a problem *during* learning if we use the cost to measure representation efficiency for this learning step. An alternative is to use a more generic L_0 norm sparseness, by simply counting the number of non-zero coefficients. It was found that by using an algorithm like Matching Pursuit, the learning algorithm could provide results similar to SPARSENET, but without the need of parametric assumptions on the prior (Perrinet, 2004; Perrinet et al., 2003; Rehn & Sommer, 2007; Smith & Lewicki, 2006). However, we observed that this class of algorithms could lead to solutions corresponding to a local minimum of the objective function: Some solutions seem as efficient as others for representing the signal but do not represent edge-like features homogeneously. In particular, during the early learning phase, some cells may learn “faster” than others. There is the need for a homeostasis mechanism that will ensure convergence of learning. The goal of this work is to study the specific role of homeostasis in learning sparse representations and to propose a homeostasis mechanism which optimizes the learning of an efficient neural representation.

To achieve this, we first formulate analytically the problem of representation efficiency in a population of sensory neurons (section 2) and define the class of Sparse Hebbian Learning (SHL) algorithms. For the particular non-parametric L_0 norm sparseness, we show that sparseness is optimal when average activity within the neural population is uniformly balanced. Based on a previous implementation, Adaptive Matching Pursuit (AMP) (Perrinet, 2004; Perrinet et al., 2003), we will define a homeostatic gain control mechanism that we will integrate in a novel SHL algorithm (section 3). Finally, we compare in section 4 this novel algorithm with AMP and the state-of-the-art SPARSENET method (Olshausen & Field, 1996). Using quantitative measures of efficiency based on constraints on the neural representation, we show the importance of the homeostasis mechanism in terms of representation efficiency. We conclude in section 5 by linking this original method with other Sparse Hebbian Learning schemes and how these may be united to improve our understanding of the emergence of edge-like simple cell receptive fields, drawing the bridge between structure (representation in a distributed network) and function (efficient coding).

2 Problem Statement

2.1 Definition of representation efficiency

In low-level sensory areas, the goal of neural computations is to generate efficient intermediate *representations* to allow efficient decision making. Classically, a representation is defined as the inversion of an internal generative model of the sensory world, that is, by inferring the sources that generated the input signal. Formally, as in Olshausen & Field (1997), we define a Linear Generative Model (LGM) for describing natural, static, grayscale images \mathbf{I} (represented by

column vectors of dimension L pixels), by setting a “dictionary” of M images (or “filters”) as the $L \times M$ matrix $\Phi = \{\Phi_i\}_{1 \leq i \leq M}$. Knowing the associated “sources” as a vector of coefficients $\mathbf{a} = \{a_i\}_{1 \leq i \leq M}$, the image is defined using matrix notation as

$$\mathbf{I} = \Phi \mathbf{a} + \mathbf{n} \quad (1)$$

where \mathbf{n} is a decorrelated gaussian additive noise image of variance σ_n^2 . The decorrelation of the noise is achieved by applying Principal Component Analysis to the raw input images, without loss of generality since this preprocessing is invertible. Generally, the dictionary Φ may be much larger than the dimension of the input space (that is, if $M \gg L$) and it is then said to be *over-complete*. However, given an over-complete dictionary, the inversion of the LGM leads to a combinatorial search and typically, there may exist many coding solutions \mathbf{a} to equation 1 for one given input \mathbf{I} . The goal of efficient coding is to find, given the dictionary Φ and for any observed signal \mathbf{I} , the “best” representation vector, that is, as close as possible to the sources that generated the signal. It is therefore necessary to define an efficiency criterion in order to choose between these different solutions.

Using the LGM, we will infer the “best” coding vector as the most probable. In particular, from the physical synthesis of natural images, we know *a priori* that image representations are sparse: They are most likely generated by a small number of features relatively to the dimension M of representation space. Similarly to Lewicki & Sejnowski (2000), this can be formalized in the probabilistic framework defined by the LGM (see equation 1), by assuming that we know the prior distribution of the coefficients a_i for natural images. The representation cost of \mathbf{a} for one given natural image is then:

$$\begin{aligned} \mathcal{C}(\mathbf{a}|\mathbf{I}, \Phi) &= -\log P(\mathbf{a}|\mathbf{I}, \Phi) \\ &= \log Z + \frac{1}{2\sigma_n^2} \|\mathbf{I} - \Phi \mathbf{a}\|^2 - \sum_i \log P(a_i|\Phi) \end{aligned} \quad (2)$$

where Z is the partition function which is independent of the coding and $\|\cdot\|$ is the L_2 norm in image space. This efficiency cost is measured in bits if the logarithm is of base 2, as we will assume without loss of generality thereafter. For any representation \mathbf{a} , the cost value corresponds to the description length (Rissanen, 1978): On the right hand side of equation 2, the second term corresponds to the information from the image which is not coded by the representation (reconstruction cost) and thus to the information that can be at best encoded using entropic coding pixel by pixel (it’s the log-likelihood in Bayesian terminology). The third term $S(\mathbf{a}|\Phi) = -\sum_i \log P(a_i|\Phi)$ is the representation or sparseness cost: It quantifies representation efficiency as the coding length of each coefficient of \mathbf{a} independently which would be achieved by entropic coding knowing the prior. In practice, the sparseness of coefficients for natural images is often defined by an *ad hoc* parameterization of the prior’s shape. For instance, the

parameterization in Olshausen & Field (1997) yields the coding cost:

$$\mathcal{C}_1(\mathbf{a}|\mathbf{I}, \Phi) = \frac{1}{2\sigma_n^2} \|\mathbf{I} - \Phi\mathbf{a}\|^2 + \beta \sum_i \log\left(1 + \frac{a_i^2}{\sigma^2}\right) \quad (3)$$

where β corresponds to the prior’s steepness and σ to its scaling (see Figure 13.2 from (Olshausen, 2002)). This choice is often favored because it results in a convex cost for which known numerical optimization methods such as conjugate gradient may be used.

A non-parametric form of sparseness cost may be defined by considering that neurons representing the vector \mathbf{a} are either active or inactive. In fact, the spiking nature of neural information demonstrates that the transition from an inactive to an active state is far more significant at the coding time scale than smooth changes of the firing rate. This is for instance perfectly illustrated by the binary nature of the neural code in the auditory cortex of rats (DeWeese et al., 2003). Binary codes also emerge as optimal neural codes for rapid signal transmission (Bethge et al., 2003; Nikitin et al., 2009). With a binary event-based code, the cost is only incremented when a new neuron gets active, regardless to the analog value. Stating that an active neuron carries a bounded amount of information of λ bits, an upper bound for the representation cost of neural activity on the receiver end is proportional to the count of active neurons, that is, to the L_0 norm:

$$\mathcal{C}_0(\mathbf{a}|\mathbf{I}, \Phi) = \frac{1}{2\sigma_n^2} \|\mathbf{I} - \Phi\mathbf{a}\|^2 + \lambda \|\mathbf{a}\|_0 \quad (4)$$

This cost is similar with information criteria such as the AIC (Akaike, 1974) or distortion rate (Mallat, 1998, p. 571). This simple non-parametric cost has the advantage of being dynamic: The number of active cells for one given signal grows in time with the number of spikes reaching the receiver (see architecture of the model in figure 1-Left). But equation 4 defines a harder cost to optimize since the hard L_0 norm sparseness leads to a non-convex optimization problem which is *NP-complete* with respect to the dimension M of the dictionary (Mallat, 1998, p. 418).

2.2 Sparse Hebbian Learning (SHL)

Given a sparse coding strategy that optimizes any representation efficiency cost as defined above, we may derive an unsupervised learning model by optimizing the dictionary Φ over natural scenes. On the one hand, the flexibility in the definition of the sparseness cost leads to a wide variety of proposed *sparse coding* solutions (for a review, see (Pece, 2002)) such as numerical optimization (Lee et al., 2007; Olshausen & Field, 1997), non-negative matrix factorization (Lee & Seung, 1999; Ranzato et al., 2007) or Matching Pursuit (Perrinet, 2004; Perrinet et al., 2003; Rehn & Sommer, 2007; Smith & Lewicki, 2006). On the other hand, these methods share the same LGM model (see equation 1) and once the sparse coding algorithm is chosen, the learning scheme is similar.

Indeed, after every coding sweep, the efficiency of the dictionary Φ may be increased with respect to equation 2. By using the online gradient descent approach given the current sparse solution, learning may be achieved using $\forall i$:

$$\Phi_i \leftarrow \Phi_i + \eta a_i (\mathbf{I} - \Phi \mathbf{a}) \quad (5)$$

where η is the learning rate. Similarly to Eq. 17 in (Olshausen & Field, 1997) or to Eq. 2 in (Smith & Lewicki, 2006), the relation is a linear “Hebbian” rule (Hebb, 1949) since it enhances the weight of neurons proportionally to the correlation between pre- and post-synaptic neurons. Note that there is no learning for non-activated coefficients. The novelty of this formulation compared to other linear Hebbian learning rule such as (Oja, 1982) is to take advantage of the sparse representation, hence the name Sparse Hebbian Learning (SHL).

SHL algorithms are unstable without homeostasis. In fact, starting with a random dictionary, the first filters to learn are more likely to correspond to salient features (Perrinet et al., 2004) and are therefore more likely to be selected again in subsequent learning steps. In SPARSENET, the homeostatic gain control is implemented by adaptively tuning the norm of the filters. This method equalizes the variance of coefficients across neurons using a geometric stochastic learning rule. The underlying heuristic is that this introduces a bias in the choice of the active coefficients. In fact, if a neuron is not selected often, the geometric homeostasis will decrease the norm of the corresponding filter, and therefore—from equation 1 and the conjugate gradient optimization—this will increase the value of the associated scalar. Finally, since the prior functions defined in equation 3 are identical for all neurons, this will increase the relative probability that the neuron is selected with a higher relative value. The parameters of this homeostatic rule have a great importance for the convergence of the global algorithm. We will now try to define a more general homeostasis mechanism derived from the optimization of representation efficiency.

2.3 Efficient cooperative homeostasis in SHL

The role of homeostasis during learning is to make sure that the distribution of neural activity is homogeneous. In fact, neurons belonging to a same neural assembly (Hebb, 1949) form a competitive network and should *a priori* carry similar information. This optimizes the coding efficiency of neural activity in terms of compression (van Hateren, 1993) and thus minimizes intrinsic noise (Srinivasan et al., 1982). Such a strategy is similar to introducing an intrinsic adaptation rule such that prior firing probability of all neurons have a similar Laplacian probability distribution (Weber & Triesch, 2008). Dually, since neural activity in the assembly actually represents the sparse coefficients, we may understand the role of homeostasis as maximizing the average representation cost $\mathcal{C}(\mathbf{a}|\Phi)$ at the time scale of learning. This is equivalent to say that homeostasis should act such that at any time, invariantly to the selectivity of features in the dictionary, the probability of selecting one feature is uniform across the dictionary.

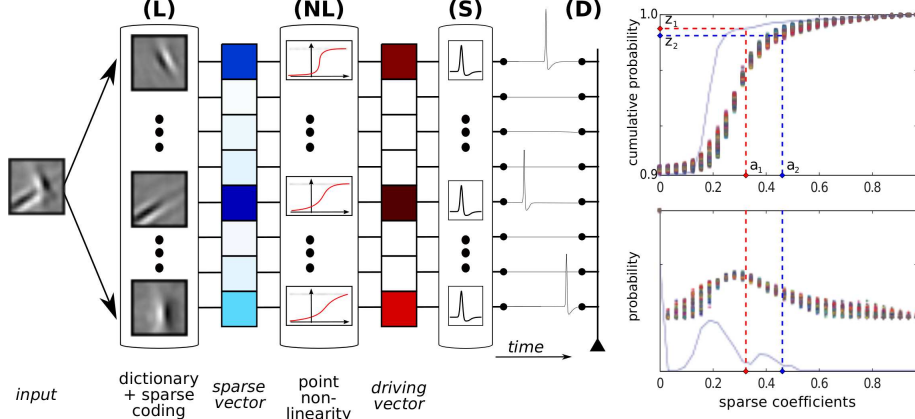


Figure 1: **Simple neural model of sparse coding and role of homeostasis.**

(Left) We define the coding model as an information channel constituted by a bundle of Linear/Non-Linear spiking neurons. **(L)** A given input image patch is coded linearly by using the dictionary of filters Φ_i and transformed by sparse coding (such as Matching Pursuit) into a sparse vector \mathbf{a} . Each coefficient is transformed into a driving coefficient in the **(NL)** layer by using a point non-linearity which **(S)** drives a generic spiking mechanism. **(D)** On the receiver end (for instance in an efferent neuron), one may then estimate the input from the neural representation pattern. This decoding is progressive, and if we assume that each spike carries a bounded amount of information, representation cost in this model increases proportionally with the number of activated neurons. (Right) However, for a given dictionary, the distribution of sparse coefficients a_i and hence the probability of a neuron's activation is in general not uniform. We show (Lower panel) the log-probability distribution function and (Upper panel) the cumulative distribution of sparse coefficients for a dictionary of edge-like filters with similar selectivity (dotted scatter) except for one filter which was randomized (continuous line). This illustrates a typical situation which may occur during learning when some components did learn less than others: Since their activity will be lower, they are less likely to be activated in the spiking mechanism and from the Hebbian rule, they are less likely to learn. When selecting an optimal sparse set for a given input, instead of comparing sparse coefficients with respect to a threshold (vertical dashed lines), it should instead be done on the significance value z_i (horizontal dashed lines): In this particular case, the less selective neuron ($a_1 < a_2$) is selected by the homeostatic cooperation ($z_1 > z_2$). The role of homeostasis during learning is that, even if the dictionary of filters is not homogeneous, the point non-linearity in (NL) modifies sparse coding in (L) such that the probability of a neuron's activation is uniform across the population.

This optimal uniformity may be achieved in all generality for any given dictionary by using point non-linearities z_i applied to the sparse coefficients: In fact, a standard method to achieve uniformity is to use an equalization of the histogram (Atick, 1992). This method may be easily derived if we know the probability distribution function dP_i of variable a_i by choosing the non-linearity as the cumulative distribution function transforming any observed variable \bar{a}_i into:

$$z_i(\bar{a}_i) = P_i(a_i \leq \bar{a}_i) = \int_{-\infty}^{\bar{a}_i} dP_i(a_i) \quad (6)$$

This is equivalent to the change of variables which transforms the sparse vector \mathbf{a} to a variable with uniform probability distribution function in $[0, 1]^M$. The transformed coefficients may thus be used as a normalized drive to the spiking mechanism of the individual neurons (see figure 1-Left). This equalization process has been observed in the neural activity of a variety of species and is, for instance, perfectly illustrated in the salamander’s retina (Laughlin, 1981). It may evolve dynamically to slowly adapt to varying changes in luminance or contrast values, such as when the light diminishes at twilight (Hosoya et al., 2005).

This novel and simple non-parametric homeostatic method is applicable to Sparse Hebbian Learning algorithms by using this transform on the sparse coefficients. Let’s imagine for instance that one filter corresponds to a feature of low selectivity while others correspond to similarly selective features: As a consequence, this filter will correspond on average to lower sparse coefficients (see figure 1-Right). However, the respective gain control function z_i will be such that all transformed coefficients have the same probability density function. Using the transformed coefficients to evaluate which neuron should be active, the homeostasis will therefore optimize the information in the representation cost defined in equation 4. We will now illustrate how it may be applied to Adaptive Matching Pursuit (Perrinet, 2004; Perrinet et al., 2003) and measure its role on the emergence of edge-like simple cell receptive fields.

3 Methods

3.1 Matching Pursuit and Adaptive Matching Pursuit

Let’s first define Adaptive Matching Pursuit. We saw that optimizing the efficiency by minimizing the L_0 norm cost leads to a combinatorial search with regard to the dimension of the dictionary. In practice, it means that for a given dictionary, finding the best sparse vector according to minimizing $\mathcal{C}_0(\mathbf{a}|\mathbf{I}, \Phi)$ (see equation 4) is hard and thus that learning an adapted dictionary is difficult. As proposed in (Perrinet et al., 2002), we may solve this problem using a greedy approach. In general, a greedy approach is applied when finding the best combination of elements is difficult to solve globally: A simpler solution is to solve the problem progressively, one element at a time.

Applied to equation 4, it corresponds to first choosing the single element $a_i\Phi_i$ that best fits the image. From the definition of the LGM, we know that for a given signal \mathbf{I} , the probability $P(\{a_i\}|\mathbf{I}, \Phi)$ corresponding to a *single* source $a_i\Phi_i$ for any i is maximal for the dictionary element i^* with maximal correlation coefficient:

$$i^* = \text{ArgMax}_i(\rho_i), \text{ with } \rho_i = \langle \frac{\mathbf{I}}{\|\mathbf{I}\|}, \frac{\Phi_i}{\|\Phi_i\|} \rangle \quad (7)$$

This formulation is slightly different from Eq. 21 in (Olshausen & Field, 1997). It should be noted that ρ_i is the L -dimensional cosine (L is the dimension of the input space) and that its absolute value is therefore bounded by 1. The value of $\text{ArcCos}(\rho_i)$ would therefore give the angle of \mathbf{I} with the pattern Φ_i and in particular, the angle (modulo 2π) would be equal to zero if and only if $\rho_i = 1$ (full correlation), π if and only if $\rho_i = -1$ (full anti-correlation) and $\pm\pi/2$ if $\rho_i = 0$ (both vectors are orthogonal, there is no correlation). The associated coefficient is the scalar projection:

$$a_{i^*} = \langle \mathbf{I}, \frac{\Phi_{i^*}}{\|\Phi_{i^*}\|^2} \rangle \quad (8)$$

Second, knowing this choice, the image can be decomposed in

$$\mathbf{I} = a_{i^*}\Phi_{i^*} + \mathbf{R} \quad (9)$$

where \mathbf{R} is the residual image. We then repeat this 2-step process on the residual (that is, with $\mathbf{I} \leftarrow \mathbf{R}$) until some stopping criterion is met.

Hence, we have a sequential algorithm which permits to reconstruct the signal using the list of choices and that we called Sparse Spike Coding (Perrinet et al., 2002). The coding part of the algorithm produces a sparse representation vector \mathbf{a} for any input image: Its L_0 norm is the number of active neurons. Note that the norm of the filters have no influence in this algorithm on the choice function nor on the cost. For simplicity and without loss of generality, we will thereafter set the norm of the filters to 1: $\forall i, \|\Phi_i\| = 1$. It is equivalent to Matching Pursuit (MP) algorithm (Mallat & Zhang, 1993) and we have proven previously that this yields an efficient algorithm for representing natural images. Using MP in the SHL scheme defined above (see section 2.2) defines Adaptive Matching Pursuit (AMP) (Perrinet, 2004; Perrinet et al., 2003) and is similar to other strategies such as (Rehn & Sommer, 2007; Smith & Lewicki, 2006). This class of SHL algorithms offers a non-parametric solution to the emergence of simple cell receptive fields, but compared to SPARSENET, the results often appear to be qualitatively non-homogeneous. Moreover, the heuristic used in SPARSENET for the homeostasis may not be used directly since in MP the choice is independent to the norm of the filter. The coding algorithm's efficiency may be improved using Optimized Orthogonal MP (Rebollo-Neira & Lowe, 2002) and be integrated in a SHL scheme (Rehn & Sommer, 2007). However, this optimization is separate with the problem that we try to tackle here by optimizing the representation at the learning time scale. Thus, we will now study how we may use cooperative homeostasis in order to optimize the overall coding efficiency of the dictionary learnt by AMP.

3.2 Competition-Optimized Matching Pursuit (COMP)

In fact, we may now include cooperative homeostasis into AMP. At the coding level, it is important to note that if we simply equalize the sparse output of the MP algorithm, transformed coefficients will indeed be uniformly distributed but the sequence of chosen filters will not be changed. However, the MP algorithm is non-linear and the choice of an element at one step may influence the rest of the choices. This sequence is therefore crucial for the representation efficiency. In order to optimize the competition of the choice step, we may instead choose at every matching step the item in the dictionary corresponding to the most significant value computed thanks to the cooperative homeostasis (see figure 1-Right). In practice, it means that we select the best match in the vector corresponding to the transformed coefficients \mathbf{z} , that is, in the vector of the residual coefficients weighted by the non-linearities defined by equation 6. This scheme thus extends the MP algorithm which we used previously by linking it to a statistical model which optimally tunes the ArgMax operator in the matching step: Over natural images, for any given dictionary —and thus independently to the selectivity of the different items from the dictionary— the choice of a neuron is statistically equally probable. Thanks to cooperative homeostasis, the efficiency of every match in MP is thus maximized, hence the name of Competition-Optimized Matching Pursuit (COMP).

Let's now explicitly describe the COMP coding algorithm step by step. Initially, given the signal \mathbf{I} , we set up for all i an internal activity vector $\bar{\mathbf{a}}$ as the linear correlation using equation 8. The output sparse vector is set initially to a zero vector: $\mathbf{a} = \mathbf{0}$. Using the internal activity $\bar{\mathbf{a}}$, the neural population will evolve dynamically in an event-based manner by repeating the two following steps. First, the “Matching” step is defined by choosing the address with the most significant activity:

$$i^* = \text{ArgMax}_i[z_i(\bar{a}_i)] \quad (10)$$

Then, we set the winning sparse coefficient at address i^* with $a_{i^*} \leftarrow \bar{a}_{i^*}$. In the second “Pursuit” step, as in MP, the information is fed-back to correlated dictionary elements by:

$$\bar{a}_i \leftarrow \bar{a}_i - a_{i^*} < \Phi_{i^*}, \Phi_i > \quad (11)$$

Note that after the update, the winning internal activity is zero: $\bar{a}_{i^*} = 0$ and that, as in MP, a neuron is selected at most once. Physiologically, as previously described, the pursuit step could be implemented by a lateral, correlation-based inhibition. The algorithm is iterated with equation 10 until some stopping criteria is reached, such as when the residual error energy is below the noise level σ_n^2 . As in MP, since the residual is orthogonal to Φ_{i^*} , the residual error energy $E = \|\mathbf{I}\|^2$ may be easily updated at every step as:

$$E \leftarrow E - a_{i^*}^2 \quad (12)$$

COMP transforms the image \mathbf{I} into the sparse vector \mathbf{a} at any precision \sqrt{E} . As in MP, the image may be reconstructed using: $\bar{\mathbf{I}} = \sum_i a_i \Phi_i$, which thus gives a

solution for equation 1. COMP differs from MP only by the “Matching” step and shares many properties with MP, such as the monotonous decrease of the error (see equation 12) or the exponential convergence of the coding. However, the decrease of E will always be faster in MP than in COMP from the constraint in the matching step.

Yet, for a given dictionary, we do not know a priori the functions z_i since they depend on the computation of the sparse coefficients. In practice, the z_i functions are initialized for all neurons to similar arbitrary cumulative distribution functions (COMP is then equivalent to the MP algorithm since choices are not affected). Since we have at most one sparse value a_i per neuron, the cumulative histogram function for each neuron for one coding sweep is $P(a_i \leq \bar{a}_i) = \delta(a_i \leq \bar{a}_i)$ where variable \bar{a}_i is the observed coefficient to be transformed and δ is the Dirac measure: $\delta(B) = 1$ if the boolean variable B is true and 0 otherwise. We evaluate equation 6 after the end of every coding using an online stochastic algorithm, $\forall i, \forall \bar{a}_i$:

$$z_i(\bar{a}_i) \leftarrow (1 - \eta_h)z_i(\bar{a}_i) + \eta_h \delta(a_i \leq \bar{a}_i) \quad (13)$$

where η_h is the homeostatic learning rate. Note that this corresponds to the empirical estimation and assumes that coefficients are stationary on a time scale of $\frac{1}{\eta_h}$ learning steps. The time scale of homeostasis should therefore in general be less than the time scale of learning. Moreover, due to the exponential convergence of MP, for any set of components, the z_i functions converge to the correct non-linear functions as defined by equation 6.

3.3 Adaptive Sparse Spike Coding (aSSC)

We may finally apply COMP to Sparse Hebbian Learning (see section 2.2). Since the efficiency is inspired by the spiking nature of neural representations, we call this algorithm adaptive Sparse Spike Coding (aSSC). From the definition of COMP, we know that whatever the dictionary, the competition between filters will be fair thanks to the cooperative homeostasis. We add no other homeostatic regulation. We normalize filters’ energy since it is a free parameter in equation 7.

In summary, the whole learning algorithm is given by the following nested loops in pseudo-code:

1. Initialize the point non-linear gain functions z_i to similar cumulative distribution functions and the components Φ_i to random points on the unit L -dimensional sphere,
2. repeat until learning converged:
 - (a) draw a signal \mathbf{I} from the database, its energy is $E = \|\mathbf{I}\|^2$,
 - (b) set sparse vector \mathbf{a} to zero, initialize $\bar{a}_i = \langle \mathbf{I}, \Phi_i \rangle$ for all i ,
 - (c) while the residual energy E is above a given threshold do:
 - i. select the best match: $i^* = \text{ArgMax}_i[z_i(\bar{a}_i)]$,




Figure 2: **Comparison of the dictionaries obtained with SPARSENET and aSSC.** We show the results of Sparse Hebbian Learning using two different sparse coding algorithms at convergence (20000 learning steps): (*Left*) conjugate gradient function (CGF) method as used in SPARSENET (Olshausen & Field, 1997) with (*Right*) COMP as used in aSSC. Filters of the same size as the image patches are presented in a matrix (separated by a black border). Note that their position in the matrix is arbitrary as in ICA.

- ii. set the sparse coefficient: $a_{i^*} = \bar{a}_{i^*}$,
 - iii. update residual coefficients: $\forall i, \bar{a}_i \leftarrow \bar{a}_i - a_{i^*} < \Phi_{i^*}, \Phi_i >$,
 - iv. update energy: $E \leftarrow E - a_{i^*}^2$.
- (d) when we have the sparse representation vector \mathbf{a} , apply $\forall i$:
- i. modify dictionary: $\Phi_i \leftarrow \Phi_i + \eta a_i (\mathbf{I} - \Phi \mathbf{a})$,
 - ii. normalize dictionary: $\Phi_i \leftarrow \Phi_i / \|\Phi_i\|$,
 - iii. update homeostasis functions: $z_i(\cdot) \leftarrow (1 - \eta_h) z_i(\cdot) + \eta_h \delta(a_i \leq \cdot)$.

4 Results on natural images

The aSSC algorithm differs from the SPARSENET algorithm by the MP sparse coding algorithm and by the cooperative homeostasis. Using natural images, we evaluate the relative contribution of these different mechanisms to the representation efficiency.

4.1 Receptive field formation

We first compare the dictionaries of filters obtained by both methods. We use a similar context and architecture as the experiments described in (Olshausen & Field, 1997) and specifically the same database of image patches as the SPARSENET

algorithm. These images are static, grayscale and whitened according to the same parameters to allow a one-to-one comparison of both algorithms. Here, we show the results for 16×16 image patches (so that $L = 256$) and the learning of $M = 324$ filters which are replicated as ON and OFF filters. Assuming this symmetry in the aSSC algorithm, we use the absolute value of the coefficient in equation 10 and equation 13¹, the rest of the algorithm being identical. Results replicate the original results of Olshausen & Field (1997) and are comparable for both methods: Dictionaries consist of edge-like filters similarly to the receptive fields of simple cells in the primary visual cortex (see figure 2). Studying the evolution of receptive fields during learning shows that they first represent any salient feature (such as sharp corners or edges), because these correspond to larger Lipschitz coefficients (Perrinet et al., 2004). If a receptive field contains multiple singularities, only the most salient remains later on during learning: Due to the competition between filters, the algorithm eliminates features that are duplicated in the dictionary. Filters which already converged to independent components will be selected sparsely and with high associated coefficients, but inducing a slower learning since corresponding error is small (see equation 5). We observe for both algorithms that when considering very long learning times, the solution is not fixed and edges may slowly drift from one orientation to another while global efficiency remains stable. This is due to the fact that there are many solutions to the same problem (note, for instance, that solutions are invariant up to a permutation of neurons' addresses). It is possible to decrease these degrees of freedom by including for instance topological links between filters (Bednar et al., 2004). Qualitatively, the main difference between both results is that filters produced by aSSC look more diverse and broad (so that they often overlap), while the filters produced by SPARSENET are more localized and thin.

We also perform robustness experiments to determine the range of learning parameters for which these algorithms converged. One advantage of aSSC is that it is based on a non-parametric sparse coding and a non-parametric homeostasis rule and is entirely described by 2 structural parameters (L and M) and 2 learning parameters (η and η_h) while parameterization of the prior and of the homeostasis for SPARSENET requires 5 more parameters to adjust (3 for the prior, 2 for the homeostasis). By observing at convergence the probability distribution function of selected filters, homeostasis in aSSC converges for a wide range of η_h values (see equation 13). Furthermore, we observe that at convergence, the z_i functions become very similar (see dotted lines in figure 1-Right) and that homeostasis does not favor the selection of any particular neuron as strongly as at the beginning of the learning. Therefore, thanks to the homeostasis, equilibrium is reached when the dictionary homogeneously represents different features in natural images, that is, when filters have similar selectivities. Finally, we observe the counter-intuitive result that non-linearities implementing cooperative homeostasis are important for the coding only *during*

¹That is, following section 3.3, step 2-c-i becomes $i^* = \text{ArgMax}_i[z_i(|\bar{a}_i|)]$, and step 2-d-iii is changed to $z_i(\cdot) \leftarrow (1 - \eta_h)z_i(\cdot) + \eta_h\delta(|a_i| \leq \cdot)$.

the learning period but that it may be ignored for the coding after convergence since at this point non-linearities are the same for all neurons.

Both dictionaries appear to be qualitatively different and for instance parameters of the emerging edges (frequency, length, width) are distributed differently. In fact, it seems that rather than the shape of each dictionary element taken individually, it is their distribution in image space that yields different efficiencies. Such an analysis of the filters' shape distribution was performed quantitatively for SPARSENET in (Lewicki & Sejnowski, 2000). The filters were fitted by Gabor functions (Jones & Palmer, 1987). A recent study compares the distribution of fitted Gabor functions' parameters between the model and receptive fields obtained from neurophysiological experiments conducted in primary visual cortex of macaques (Rehn & Sommer, 2007). It has shown that their SHL model based on Optimized Orthogonal MP better matches to physiological observations than SPARSENET. However, there is no theoretical basis for the fact that receptive fields' shape should be well fitted by Gabor functions (Saito, 2001) and the variety of shapes observed in biological systems may for instance reflect adaptive regulation mechanisms when reaching different optimal sparseness levels (Assisi et al., 2007). Moreover, even though this type of quantitative method is certainly necessary, it is not sufficient to understand the role of each individual mechanism in the emergence of edge-like receptive fields. To assess the relative role of coding and homeostasis in SHL, we rather compare these different dictionaries quantitatively in terms of representation efficiency.

4.2 Coding efficiency in SHL

To address this issue, we first compare the quality of both dictionaries (from SPARSENET and aSSC) by computing the mean efficiency of their respective *coding* algorithms (respectively CGF and COMP). Using 10^5 image patches drawn from the natural image database, we perform the progressive coding of each image using both sparse coding methods. When plotting the probability distribution function of the sparse coefficients, one observes that distributions fit well the bivariate model introduced in (Olshausen & Millman, 2000) where a sub-set of the coefficients are null (see figure 3-Left). Log-probability distributions of non-zero coefficients is quadratic with the initial random dictionaries. At convergence, non-zero coefficients fit well to a Laplacian probability distribution function. Measuring mean kurtosis of resulting sparse vectors proves to be very sensitive and a poor indicator of global efficiency, in particular at the beginning of the coding, when many coefficients are still strictly zero. In general, COMP provides a sparser final distribution. Dually, plotting the decrease of the sorted coefficients as a function of their rank shows that coefficients for COMP are first higher and then decrease more quickly, due to the link between the z_i functions and the function of sorted coefficients (see equation 6). As a consequence, a Laplacian bivariate model for the distribution of sparse coefficient emerge from the statistics of natural images. The advantage of aSSC is that this emergence is not dependent of a parametric model of the prior.

In a second analysis, we compare the efficiency of both methods while varying




Figure 3: **Coding efficiency of SPARSENET versus aSSC.** We evaluate the quality of both learning schemes by comparing coding efficiency of their respective coding algorithms, that is CGF and COMP, with the respective dictionary that was learnt (see figure 2). (*Left*) We show the probability distribution function of sparse coefficients obtained by both methods with random dictionaries (respectively 'SN-init' and 'aSSC-init') and with the dictionaries obtained after convergence of respective learning schemes (respectively 'SN' and 'aSSC'). At convergence, sparse coefficients are more sparsely distributed than initially, with more kurtotic probability distribution functions for aSSC in both cases. (*Right*) We plot the average residual error (L_2 norm) as a function of the relative number of active (non-zero) coefficients. This provides a measure of the coding efficiency for each dictionary over the set of image patches (error bars are scaled to one standard deviation). The L_0 norm is equal to the coding step in COMP. Best results are those providing a lower error for a given sparsity (better compression) or a lower sparseness for the same error (Occam's razor). We observe similar coding results in aSSC despite its non-parametric definition. This result is also true when using the two different dictionaries with the same OOMP sparse coding algorithm: The dictionaries still have similar coding efficiencies.

the number of active coefficients (the L_0 norm). We perform this in COMP by simply measuring the residual error (L_2 norm) with respect to the coding step. To compare this method with the conjugate gradient method, we use a 2-pass sparse coding: A first pass identifies best neurons for a fixed number of active coefficients, while a second pass optimizes the coefficients for this set of “active” vectors. This method was also used in (Rehn & Sommer, 2007) and proved to be fair when comparing both algorithms. We observe in a robust manner that the greedy solution to the hard problem (that is, COMP) is as efficient as conjugate gradient as used in SPARSENET (see figure 3, Right). We also observe that aSSC is also slightly more efficient for the cost defined in equation 3, a result which may reflect the fact that the L_0 norm defines a stronger sparseness constraint than the convex cost. Moreover, we compare the coding efficiency of both dictionaries using Optimized Orthogonal MP. Results show that OOMP provides a slight coding improvement, but also confirms that both dictionaries are of similar coding efficiency, independently of their respective coding algorithm.

These results prove that, without the need of a parameterization of the prior, coding in aSSC is as efficient as SPARSENET. In addition, there are a number of other advantages offered by this approach. First, COMP simply uses a feed-forward pass with lateral interactions, while conjugate gradient is implemented as the fixed point of a recurrent network (see Figure 13.2 from (Olshausen, 2002)). Moreover, we have already seen that aSSC is a non-parametric method which is controlled by fewer parameters. Therefore, applying a “higher-level” Occam razor confirms that for a similar overall coding efficiency, aSSC is better since it is of lower *structural* complexity². Finally, in SPARSENET and in algorithms defined in (Lewicki & Sejnowski, 2000; Rehn & Sommer, 2007; Smith & Lewicki, 2006), representation is analog without explicitly defining a quantization. This is not the case in the aSSC algorithm where cooperative homeostasis introduces a regularity in the distribution of sparse coefficients.

4.3 Role of homeostasis in representation efficiency

In the context of an information channel such as implemented by a neural assembly, one should rather use the coefficients that could be decoded from the neural signal in order to define the reconstruction cost (see figure 1, Left). As was described in section 2.1, knowing a dictionary Φ , it is indeed more correct to consider the overall average coding and decoding cost over image patches $\mathcal{C}(\hat{\mathbf{a}}|\mathbf{I}, \Phi)$ (see equation 2), where $\hat{\mathbf{a}}$ corresponds to the analog vector of coefficients inferred from the neural representation. The overall transmission error may be described as the sum of the reconstruction and the quantization error. This last error will increase both with inter-trial variability but also with the

²A quantitative measure of the structural complexity for the different methods is given by the minimal length of a code that would implement them, this length being defined as the number of characters of the code implementing the algorithm. It would therefore depend on the machine on which it is implemented, and there is, of course, a clear advantage of aSSC on parallel architectures.

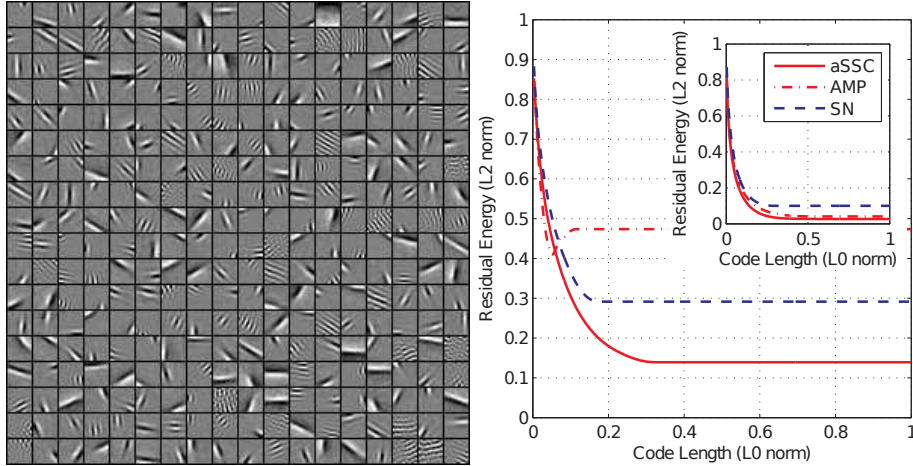


Figure 4: **Cooperative homeostasis implements efficient quantization.** (*Left*) When switching off the cooperative homeostasis during learning, the corresponding Sparse Hebbian Learning algorithm, Adaptive Matching Pursuit (AMP), converges to a set of filters that contains some less localized filters and some high-frequency Gabor functions that correspond to more “textural” features (Perrinet et al., 2003). One may wonder if these filters are inefficient and capturing noise or if they rather correspond to independent features of natural images in the LGM model. (*Right, Inset*) In fact, when plotting residual energy as a function of L_0 norm sparseness with the MP algorithm (as plotted in figure 3, Right), the AMP dictionary gives a slightly worse result than aSSC. (*Right*) Moreover, one should consider representation efficiency as the overall coding and decoding algorithm. We compare the efficiency for these dictionaries thanks to same coding method (SSC) and the same decoding method (using rank quantized coefficients). Representation length for this decoding method is proportional to the L_0 norm with $\lambda = \frac{\log(M)}{L} \approx 0.032$ bits per coefficient and per pixel as defined in equation 4. We observe that the dictionary obtained by aSSC is more efficient than the one obtained by AMP while the dictionary obtained with SPARSENET (SN) gives an intermediate result thanks to the geometric homeostasis: Introducing cooperative homeostasis globally improves neural representation.

non homogeneity of the represented features. It is however difficult to evaluate a decoding scheme in most sparse coding algorithms since this problem is generally not addressed. Our objective when defining $\mathcal{C}_0(\hat{\mathbf{a}}|\mathbf{I}, \Phi)$ (see equation 4) was to define sparseness as it may be represented by spiking neural representations. Using a decoding algorithm on such a representation will help us to quantify overall coding efficiency.

An effective decoding algorithm is to estimate the analog values of the sparse vector (and thus reconstruct the signal) from the order of neurons' activation in the sparse vector (Perrinet, 2007, Section 2.2). In fact, knowing the address of the fiber i^0 corresponding to the maximal value, we may infer that it has been produced by an analog value on the emitter side in the highest quantile of the probability distribution function of a_{i^0} . We may therefore decode the corresponding value with the best estimate which is given as the average maximum sparse coefficient for this neuron by inverting z_{i^0} (see equation 6)³: $a_{i^0} = z_{i^0}^{-1}(1)$. This is also true for the following coefficients. We write as $\frac{r}{M}$ the relative rank of the r^{th} and o the order function which gives the address of the winning neuron at rank r . Since $z_{o(r)} = 1 - \frac{r}{M} = z_{o(r)}(a_{o(r)})$, we can reconstruct the corresponding value as

$$\hat{a}_{o(r)} = z_{o(r)}^{-1}\left(1 - \frac{r}{M}\right) \quad (14)$$

Physiologically, equation 14 could be implemented using interneurons which would “count” the number of received spikes and by modulating efficiency of synaptic events on receiver efferent neurons—for instance with shunting inhibition (Delorme & Thorpe, 2003). Recent findings show that this type of code may be used in cortical in vitro recurrent networks (Shahaf et al., 2008). This corresponds to a generalized rank coding scheme. However this quantization does not require that neural information explicitly carries rank information. In fact, this scheme is rather general and is analogous to scalar quantization using the modulation function z_i^{-1} as a Look-Up-Table. It is very likely that fine temporal information such as inter-spike intervals also play a role in neural information transmission. As in other decoding schemes, the quantization error directly depends on the variability of the modulation functions across trials (Perrinet et al., 2004). This scheme thus rather shows a representative behavior for the retrieval of information from spiking neural activity.

To evaluate the specific role of cooperative homeostasis, we compare previous dictionaries (see figure 2) with the one obtained by Adaptive Matching Pursuit (AMP). In fact, SPARSENET and aSSC differ at the level of the homeostasis but also for the sparse coding. The only difference between aSSC and AMP is the introduction of cooperative homeostasis. To obtain the solution to AMP, we use the same sparse coding algorithm but switch off the cooperative homeostasis during learning ($\eta_h = 0$ in equation 13). We observe at convergence that the dictionary corresponds qualitatively to features which are different from aSSC and SPARSENET (see figure 4, Left). In particular, we observe the emergence of

³Mathematically, the z_i are not always *strictly* increasing and we state here that $z_i^{-1}(z)$ is defined in a unique way as the average value of the coefficients a_i such that $z_i(a_i) = z$.

Gabor functions with broader width which better match textures. These filters correspond to lower Lipschitz coefficients (Perrinet et al., 2004), and because of their lower saliency, these textural filters are more likely to be selected with lower correlation coefficients. They fit more to Fourier filters that are obtained using Principal Component Analysis (Fyfe & Baddeley, 1995) and are still optimal to code arbitrary image patches such as noise (Zhaoping, 2006). When we plot L_2 norm with respect to L_0 norm for the different dictionaries with the same MP coding algorithm averaged over a set of 10^5 image patches from natural scenes (see figure 4, Right Inset), the resulting dictionary from AMP is less efficient than those obtained with aSSC and SPARSENET. This is not an expected behavior since COMP is more constrained than MP (MP is the “greediest” solution) and using both methods with a similar dictionary would necessarily give an advantage to MP: the AMP thus reached a local minima of the coding cost. To understand why, recall that in the aSSC algorithm, the cooperative homeostasis constraint, by its definition in equation 6, plays the role of a gain control and that the point non-linearity from equation 10 ensures that all filters are selected equally. Compared to AMP, textured elements are “boosted” during learning relative to a more generic salient edge component and are thus more likely to evolve (see figure 1, Right). This explains why they would end up being less probable and that at convergence there are no textured filters in the dictionary obtained with aSSC.

Finally, we test quantitatively representation efficiency of these different dictionaries with the same quantization scheme. At the decoding level, we compute in all cases the modulation functions as defined in equation 14 on a set of 10^5 image patches from natural scenes. Since addresses’ choices may be generated by any of the M neurons, the representation cost is defined as $\lambda = \log(M)$ bits per chosen address (see equation 4). Then, when using the quantization (see equation 14), the AMP approach displays a larger variability reflecting the lack of homogeneity of the features represented by the dictionary: There is a much larger reconstruction error and a slower decrease of error’s energy (see figure 4, Right). The aSSC on the contrary is adapted to quantization thanks to the cooperative homeostasis and consequently yields a more regular decrease of coefficients as a function of rank, that is, a lower quantization error. The dictionary obtained with the SPARSENET algorithm yields an intermediate result. This shows that the heuristic implementing the homeostasis in this algorithm regulates relatively well the choices of the elements during the learning. It also explains why the three parameters of the homeostasis algorithm had to be properly tuned to fit the dynamics of the heuristics. Results therefore show that homeostasis optimizes the efficiency of the neural representation during learning and that the cooperative homeostasis provides a simple and effective optimization scheme.

5 Discussion

We have shown in this paper that homeostasis plays an essential role in Sparse Hebbian Learning (SHL) schemes and thus on our understanding of the emergence of simple cell receptive fields. First, using statistical inference and information theory, we have proposed a quantitative cost for the coding efficiency based on a non-parametric model using the number of active neurons, that is, the L_0 norm of the representation vector. This allowed to design a cooperative homeostasis rule based on neurophysiological observations (Laughlin, 1981). This rule optimizes the competition between neurons by simply constraining the choice of every selection of an active neuron to be equiprobable. This homeostasis defined a new sparse coding algorithm, COMP, and a new SHL scheme, aSSC. Then, we have confirmed that the aSSC scheme provides an efficient model for the formation of simple cell receptive fields, similarly to other approaches. The sparse coding algorithms in these schemes are variants of conjugate gradient or of Matching Pursuit. They are based on correlation-based inhibition since this is necessary to remove redundancies from the linear representation. This is consistent with the observation that lateral interactions are necessary for the formation of elongated receptive fields (Bolz & Gilbert, 1989). With a correct tuning of parameters, all schemes show the emergence of edge-like filters. The specific coding algorithm used to obtain this sparseness appears to be of secondary importance as long as it is adapted to the data and yields sufficiently efficient sparse representation vectors. However, resulting dictionaries vary qualitatively among these schemes and it was unclear which algorithm is the most efficient and what was the individual role of the different mechanisms that constitute SHL schemes. At the learning level, we have shown that the homeostasis mechanism had a great influence on the qualitative distribution of learned filters. In particular, using the comparison of coding and decoding efficiency of aSSC with and without this specific homeostasis, we have proven that cooperative homeostasis optimized overall representation efficiency. This efficiency is comparable with that of SPARSENET, but with the advantage that our unsupervised learning model is non-parametric and does not need to be properly tuned.

This work might be advantageously applied to signal processing problems. First, we saw that optimizing the representation cost maximizes the independence between features and is related to the goal of ICA. Since we have built a solution to the LGM inverse problem that is more efficient than standard methods such as the SPARSENET algorithm, it is thus a good candidate solution to Blind Source Separation problems. Second, at the coding level, we optimized in the COMP algorithm the efficiency of Matching Pursuit by including an adaptive cooperative homeostasis mechanism. We proved that for a given compression level, image patches are more efficiently coded than in the Matching Pursuit algorithm. Since we have shown previously that MP compares favorably with compression methods such as JPEG with a fixed log-Gabor filter dictionary (Fischer et al., 2007), we can predict that COMP should provide promising results for image representation. An advantage over other sparse

coding schemes is that it provides a progressive dynamical result while the conjugate gradient method has to be recomputed for every different number of coefficients. The most relevant information is propagated first and progressive reconstruction may be interrupted at any time. Finally, a main advantage of this type of neuromorphic algorithm is that it uses a simple set of operations: computing the correlation, applying the point non-linearity from a Look-Up Table, choosing the ArgMax, doing a subtraction, retrieving a value from a Look-Up-Table. In particular, the complexity of these operations, such as the ArgMax operator, would in theory not depend on the dimension of the system in parallel machines and the transfer of this technology to neuromorphic hardware such as aVLSIs (Brüderle et al., 2009; Schemmel et al., 2006) will provide a supra-linear gain of performance.

In this paper, we focused on transient input signals and of relatively abstract neurons. This choice was made to highlight the powerful function of the parallel and temporal competition between neurons in contrast to traditional analog and sequential strategies using analog spike frequency representations. This strategy allowed to compare the proposed learning scheme with state-of-the-art algorithms. One obvious extension to the algorithm is to implement learning with more realistic inputs. In fact, sparseness in image patches is only local while it is also spatial and temporal in whole-field natural scenes: For instance, it is highly probable in whole natural images that large parts of the space — such as the sky — are flat and contain no information. Our results should be thus taken as a lower bound for the efficiency of aSSC in natural scenes. This also suggests the extension to representations with some built-in invariances, such as translation and scaling. A gaussian pyramid, for instance, provides a multi-scale representation where the set of learned filters would become a dictionary of mother wavelets (Perrinet, 2007, Section 3.3.4). Such an extension leads to a fundamental question: How does representation efficiency evolves with the number M of elements in the dictionary, that is, with the complexity of the representation? In fact, when increasing the over-completeness in aSSC, one observes the emergence of different classes of edge filters: at first different positions, then different orientations of edges, followed by different frequencies and so on and so forth. This specific order indicates the existence of an underlying hierarchy for the synthesis of natural scenes. This hierarchy seems to correspond to the level of importance of the different transformations that are learned by the system, respectively translation, rotation and scaling. Exploring the efficiency results for different dimensions of the dictionary in aSSC will thus give a quantitative evaluation of the optimal complexity of the model needed to describe images in terms of a trade-off between accuracy and generality. But it may also provide a model for the clustering of the low-level visual system into different areas, such as the emergence of position-independent representations in the ventral visual pathway versus motion-selective neurons in the dorsal visual pathway.

Acknowledgments

This work was supported by a grant from the French Research Council (ANR “NatStats”) and by EC IP project FP6-015879, “FACETS”. The author thanks the team at the Redwood Neuroscience Institute for stimulating discussions and particularly Jeff Hawkins, Bruno Olshausen, Fritz Sommer, Tony Bell, Dileep George, Kilian Koepsell and Matthias Bethge. Special thanks to Jo Hausmann, Guillaume Masson, Nicole Voges, Willie Smit and Artemis Kosta for essential comments on this work. We thank the anonymous referees for their helpful comments on the manuscript. Special thanks to Bruno Olshausen, Laura Rebollo-Neira, Gabriel Peyré, Martin Rehn and Fritz Sommer for providing the source code for their experiments.

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19, 716–23.
- Assisi, C., Stopfer, M., Laurent, G., & Bazhenov, M. (2007). Adaptive regulation of sparseness by feedforward inhibition. *Nature Neuroscience*, 10(9), 1176–84.
- Atick, J. J. (1992). Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems*, 3(2), 213–52.
- Barlow, H. B. (2001). Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12, 241–25.
- Baudot, P., Levy, M., Monier, C., Chavane, F., René, A., Huguet, N., Marre, O., Pananceau, M., Kopysova, I., & Fregnac, Y. (2004). Time-coding, low noise Vm attractors, and trial-by-trial spiking reproducibility during natural scene viewing in V1 cortex. In *Society for Neuroscience Abstracts: 34th Annual Meeting of the Society for Neuroscience, San Diego, USA* (Eds.), (pp. 948–12).
- Bednar, J. A., Kelkar, A., & Miikkulainen, R. (2004). Scaling self-organizing maps to model large cortical networks. *Neuroinformatics*, 2(3), 275–302.
- Bethge, M., Rotermund, D., & Pawelzik, K. (2003). Second order phase transition in neural rate coding: Binary encoding is optimal for rapid signal transmission. *Physical Review Letters*, 90(8), 088104.
- Bolz, J., & Gilbert, C. D. (1989). The role of horizontal connections in generating long receptive fields in the cat visual cortex. *European Journal of Neuroscience*, 1(3), 263–8.
- Brüderle, D., Müller, E., Davison, A., Müller, E., Schemmel, J., & Meier, K. (2009). Establishing a novel modeling tool: a python-based interface for a neuromorphic hardware system. *Frontiers in Neuroinformatics*, 3, 17.

- Chapman, B., & Stryker, M. P. (1992). Origin of orientation tuning in the visual cortex. *Current Opinion in Neurobiology*, 2(4), 498–501.
- Delorme, A., & Thorpe, S. J. (2003). Early cortical orientation selectivity: How fast shunting inhibition decodes the order of spike latencies. *Journal of Computational Neuroscience*, 15, 357–65.
- DeWeese, M. R., Wehr, M., & Zador, A. M. (2003). Binary coding in auditory cortex. *Journal of Neuroscience*, 23(21).
- Doi, E., Balcan, D. C., & Lewicki, M. S. (2007). Robust coding over noisy overcomplete channels. *IEEE Transactions in Image Processing*, 16(2), 442–52.
- Field, D. J. (1994). What is the goal of sensory coding? *Neural Computation*, 6(4), 559–601.
- Fischer, S., Redondo, R., Perrinet, L., & Cristóbal, G. (2007). Sparse approximation of images inspired from the functional architecture of the primary visual areas. *EURASIP Journal on Advances in Signal Processing*, 2007(1), 122.
- Fyfe, C., & Baddeley, R. J. (1995). Finding compact and sparse- distributed representations of visual images. *Network: Computation in Neural Systems*, 6, 333–44.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley.
- Hosoya, T., Baccus, S. A., & Meister, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436(7047), 71–7.
- Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6), 1233–58.
- Laughlin, S. B. (1981). A simple coding procedure enhances a neuron’s information capacity. *Zeitung für Naturforschung*, 9–10(36), 910–2.
- Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401, 788–91.
- Lee, H., Battle, A., Raina, R., & Ng, A. (2007). Efficient sparse coding algorithms. In B. Schölkopf, J. Platt, & T. Hoffman (Eds.) *Advances in Neural Information Processing Systems 19*, (pp. 801–808). Cambridge, MA: MIT Press.
- Lewicki, M. S., & Sejnowski, T. J. (2000). Learning overcomplete representations. *Neural Computation*, 12(2), 337–65.

- Mallat, S. (1998). *A wavelet tour of signal processing*. Academic Press, second edition.
- Mallat, S., & Zhang, Z. (1993). Matching Pursuit with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12), 3397–3414.
- Nikitin, A. P., Stocks, N. G., Morse, R. P., & McDonnell, M. D. (2009). Neural population coding is optimized by discrete tuning curves. *Physical Review Letters*, 103(13), 138101.
- Oja, E. (1982). A Simplified Neuron Model as a Principal Component Analyzer. *Journal of Mathematical biology*, 15, 267–73.
- Olshausen, B. A. (2002). Sparse codes and spikes. In R. P. N. Rao, B. A. Olshausen, & M. S. Lewicki (Eds.) *Probabilistic Models of the Brain: Perception and Neural Function*, chap. Sparse Codes and Spikes, (pp. 257–72). MIT Press.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583), 607–9.
- Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Research*, 37, 3311–25.
- Olshausen, B. A., & Millman, K. J. (2000). Learning sparse codes with a mixture-of-gaussians prior. In M. I. Jordan, M. J. Kearns, & S. A. Solla (Eds.) *Advances in neural information processing systems*, vol. 12, (pp. 887–93). The MIT Press, Cambridge, MA.
- Pece, A. E. C. (2002). The problem of sparse image coding. *Journal of Mathematical Imaging and Vision*, 17, 89–108.
- Perrinet, L. (2004). Finding Independent Components using spikes : a natural result of hebbian learning in a sparse spike coding scheme. *Natural Computing*, 3(2), 159–75.
- Perrinet, L. (2007). Dynamical neural networks: modeling low-level vision at short latencies. In *Topics in Dynamical Neural Networks: From Large Scale Neural Networks to Motor Control and Vision*, vol. 142 of *The European Physical Journal (Special Topics)*, (pp. 163–225). Springer Verlag (Berlin / Heidelberg).
- Perrinet, L., Samuelides, M., & Thorpe, S. J. (2002). Sparse spike coding in an asynchronous feed-forward multi-layer neural network using Matching Pursuit. *Neurocomputing*, 57C, 125–34.
- Perrinet, L., Samuelides, M., & Thorpe, S. J. (2003). Emergence of filters from natural scenes in a sparse spike coding scheme. *Neurocomputing*, 58–60(C), 821–6.

- Perrinet, L., Samuelides, M., & Thorpe, S. J. (2004). Coding static natural images using spiking event times: do neurons cooperate? *IEEE Transactions on Neural Networks*, 15(5), 1164–75. Special issue on 'Temporal Coding for Neural Information Processing'.
- Ranzato, M. A., Poultney, C. S., Chopra, S., & LeCun, Y. (2007). Efficient learning of sparse overcomplete representations with an energy-based model. In B. Schölkopf, J. Platt, & T. Hoffman (Eds.) *Advances in neural information processing systems*, vol. 19, (pp. 1137–44). Cambridge, MA: The MIT Press.
- Rebollo-Neira, L., & Lowe, D. (2002). Optimized orthogonal matching pursuit approach. *IEEE Signal Processing Letters*, 9(4), 137–40.
- Rehn, M., & Sommer, F. T. (2007). A model that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields. *Journal of Computational Neuroscience*, 22(2), 135–46.
- Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, 14, 465–71.
- Saito, N. (2001). The generalized spike process, sparsity, and statistical independence. In D. N. Rockmore, & D. M. Healy (Eds.) *Modern Signal Processing*, (p. 317). Cambridge University Press.
- Schemmel, J., Gruebl, A., Meier, K., & Mueller, E. (2006). Implementing synaptic plasticity in a VLSI spiking neural network model. In IEEE Press (Ed.) *Proceedings of the 2006 International Joint Conference on Neural Networks (IJCNN'06)*.
- Shahaf, G., Eytan, D., Gal, A., Kermany, E., Lyakhov, V., Zrenner, C., & Marom, S. (2008). Order-based representation in random networks of cortical neurons. *PLoS Computational Biology*, 4(11), e1000228+.
- Smith, E. C., & Lewicki, M. S. (2006). Efficient auditory coding. *Nature*, 439(7079), 978–82.
- Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: A fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 216(1205), 427–59.
- van Hateren, J. H. (1993). Spatiotemporal contrast sensitivity of early vision. *Vision Research*, 33, 257–67.
- Vinje, W. E., & Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287, 1273–1276.
- Weber, C., & Triesch, J. (2008). A sparse generative model of V1 simple cells with intrinsic plasticity. *Neural Computation*, 20(5), 1261–84.

- Zhaoping, L. (2006). Theoretical understanding of the early visual processes by data compression and data selection. *Network: Computation in Neural Systems*, 17(4), 301–34.
- Zibulevsky, M., & Pearlmutter, B. A. (2001). Blind Source Separation by sparse decomposition. *Neural Computation*, 13(4), 863–82.